(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 750 291 A1

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
27.12.1996 Bulletin 1996/52

(51) Int. Cl.[6]: G10L 3/00

(21) Application number: 93203407.7

(22) Date of filing: 29.05.1987

(84) Designated Contracting States:
AT BE CH DE ES FR GB IT LI NL SE

(30) Priority: 02.06.1986 GB 8613327

(62) Application number of the earlier application in accordance with Art. 76 EPC: 87304793.0

(71) Applicant: BRITISH TELECOMMUNICATIONS public limited company
London EC1A 7AJ (GB)

(72) Inventor: Forse, Nicholas John Arnold
Stoke Park, Ipswich, Suffolk (GB)

(74) Representative: Roberts, Simon Christopher et al
BT Group Legal Services,
Intellectual Property Department,
8th Floor, Holborn Centre
120 Holborn
London EC1N 2TE (GB)

Remarks:
This application was filed on 03.12.1993 as a divisional application to the application mentioned under INID code 62.

(54) Speech processor

(57) In a speech processor such as a speech recogniser, the problem of detecting the beginning and end of speech or a word accurately, to enable the creation of a speech or a word template which consistently matches stored speech or word templates is solved by characterising background noise and forming a background noise template, setting a speech threshold above which speech is detected and stored, and subtracting the background noise template from the stored speech to form a speech template.

Fig.1.

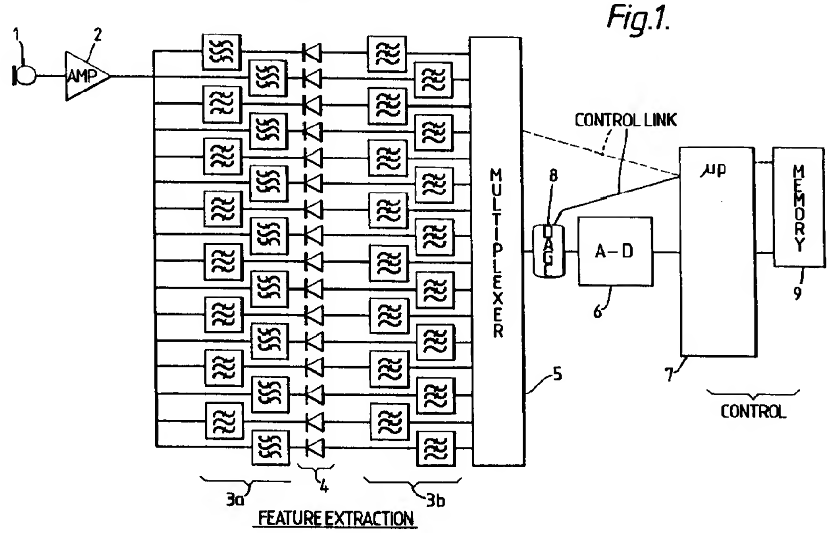

EP 0 750 291 A1

Description

This invention relates to speech processors, and in particular to speech recognisers which are able to detect the beginning and end of input speech.

Automatic speech recognisers work by comparing features extracted from audible speech signals. Features extracted from the speech to be recognised are compared with stored features extracted from a known utterance.

For accurate recognition it is important that the features extracted from the same word or sound then spoken at different times are sufficiently similar. However, the large dynamic range of speech makes this difficult to achieve, particularly in areas such as hands-free telephony where the sound level received by the microphone can vary over a wide range. In order to compensate for this speech level variation, most speech recognisers use some form of automatic gain control (AGC).

The AGC circuit controls the gain to ensure that the average signal level used by the feature extractor is as near constant as possible over a given time period. Hence quiet speech utterances are given greater gain than loud utterances. This form of AGC performs well when continuous speech is the input signal since after a period of time, the circuit gain will optimise the signal level to give consistent feature extraction. However, in the absence of speech, the gain of the AGC circuit will increase to a level determined by the background noise, so that at the onset of a speech utterance the gain of the AGC circuit will be set too high. During the utterance the gain of the circuit is automatically reduced, the speed of the gain change being determined by the 'attack' time of the AGC. The start of the utterance is thus subjected to a much greater gain and any features extracted will have a much greater energy content than similar features extracted later, when the gain has been reduced.

This distortion effect is dependent on the input signal level; the higher the speech level the larger is the distortion. Hence the first few features extracted will not correspond to the notionally similar stored features, and this can often result in poor recognition performance.

Recognition performance, as well as depending on how the AGC deals with noise level, depends on the speech recogniser's ability to detect the beginning and end of the speech or word with a high degree of accuracy, where the speech or word will become the subject of a speech template.

The present invention seeks to provide a solution to the problem of detecting the beginning and end of the speech or word with a high degree of accuracy.

According to a first aspect, the present invention provides a method of processing speech in which the start and end points of a speech sample in an input signal are determined by establishing an initial threshold based on a measure of the noise level in said input signal, the initial threshold being used to establish an initial stored speech sample, which initial speech sample is then further processed using a further threshold level which is at a predetermined level beneath a maximum level of said initial speech sample, said further threshold level being used to determine the start and end points.

According to a second aspect, the present invention provides a method of detecting speech comprising the steps of:

determining a speech energy threshold for an input channel;
forming a background noise template for the input channel;
sampling the input channel and storing the samples into a first store;
detecting speech by determining when a sample exceeds the speech threshold, and storing the sample and subsequent samples into a second store;
detecting the end of the speech by determining when a predefined number of the subsequent samples drop below the threshold;
forming a first speech template by augmenting the samples from the second store with a predefined number of samples from the first store; and
forming a second speech template by subtracting the background noise template from the first speech template.

In a method of detecting speech according to the second aspect of the present invention, it is necessary to characterise the background noise before speech is detected, so that the AGC level can be set to cope with the maximum noise level. A template representative of the background noise is formed and the speech threshold is set to be a specified amount above the maximum background noise level, so that speech is detected as soon as the speech threshold is exceeded. Also, the end of the speech is detected by recognising when the noise level drops below a threshold for a specified interval. Throughout speech detection, the AGC compensates for increasing or decreasing noise levels so that, ultimately, a speech template can be provided which has constant gain.

Embodiments of the invention will be further described and explained by way of example only, with reference to the accompanying drawing, in which;

Figure 1 is a schematic diagram of a speech recogniser according to the present invention.

Throughout this patent application the invention is described with reference to a speech recogniser utilising tem-

plate-matching, but as those skilled in the art will be aware, the invention is equally applicable to any of the conventional types of speech recogniser, including those using stochastic modelling, Markov chains, dynamic-timewarping and phoneme-recognition.

Speech recognition is based on comparing energy contours from a number (generally 8 to 16) of filter channels. While speech is present, the energy spectrum from each filter channel is digitised with an Analogue to Digital (A-D) converter to produce a template which is stored in a memory.

The initial stage of recognition is known as 'training' and consists of producing the reference templates by speaking to the recogniser the words which are to be recognised. Once reference templates have been made for the words to be recognised, recognition of speech can be attempted.

When the recogniser is exposed to an utterance, it produces a test template which can be compared with the reference templates in the memory to find the closest match.

The fundamental elements of the speech recogniser according to the present invention are shown in Figure 1. Voice signals received by the microphone 1 and amplified by amplifier 2 are passed to a filter bank 3a. In the filter bank the voice signals are filtered into a plurality (in this case 16) of frequency bands, and the signals are rectified by rectifier 4. The filtered and rectified signals are smoothed by low pass filters 3b and then sequentially sampled by a multiplexer 5 which feeds the resultant single channel signal to the DAGC circuit 8 which in turn feeds an Analogue to Digital converter 6 from which the digitised signal stream is passed to the controlling microprocessor 7.

The multiplexer addresses each filter channel for 20 microseconds before addressing the next one. At the end of each 10 millisecond time slot, each channel's sampled energy for that period is stored. The templates, which are produced during training or recognition, consist of up to 100 time slot samples for each filter channel.

The digital AGC operates in the following way. Each time the multiplexer addresses a filter channel, the microprocessor assesses the channel's energy level to determine whether the A-D convertor has been overloaded and hence that the gain is too high. When the microprocessor determines that the gain is too high it decrements the AGC's gain by 1 step, which corresponds to a reduction in gain of 1.5dB, and looks again at the channel's energy level. The multiplexer does not cycle to the next channel until the microprocessor has determined that the gain has been reduced sufficiently to prevent overloading of the A-D converter. When the multiplexer does cycle to the next filter channel, the gain of the AGC circuit is held at the new low level unless that level results in the overloading of the A-D converter with the new channel's energy level, in which case the gain is incremented down as previously described. When the multiplexer has addressed the final filter channel, the microprocessor normalises the energy levels of all the channels by setting their gain coefficients (which have been stored together with the energy level information in memory associated with the microprocessor) to the new minimum established by the microprocessor. In this way a consistent set of features are extracted independent of the initial input signal gain and any changes in the gain during formation of the template.

The speech recogniser is required to detect the beginning and end of the speech or word with a high degree of accuracy. The speech recogniser according to the present invention uses the following technique:

A
The energy level of the background noise is measured and stored for 32 time slots (at 10 milliseconds a sample) while simultaneously adjusting (reducing) the gains of the AGC circuit as described above to cope with the maximum noise energy.

B
The maximum energy sample is found by adding all the filter values for each time slot, dividing by 16 (the number of filter channels) and multiplying by a gain factor corresponding to the gain of the DAGC circuit, and then comparing each time slot to find the maximum.

C
The threshold which needs to be exceeded before speech is deemed to be present is set to be equal to 1.5 times the maximum noise energy determined in Step B.

D
The average noise energy for each filter channel is found and stored (for each channel it is the sum of energies over all 32 time slots, divided by 32) to establish a noise template.

E
Thereafter, the filter bank is scanned every 10 milliseconds and the data is stored in a temporary cyclic store, of 100 time samples, until the average filter energy exceeds the noise/speech threshold calculated in C.

F
If the noise/speech threshold is not exceeded after 32 samples, a check is performed to ensure that the gain of the

DAGC circuit is not set too low. This is done by looking at the maximum filter channel value stored in those 32 time slots. If that maximum level is 1.5dB or more below the maximum acceptable input level for the A-D converter, the gain of the AGC is incremented by 1 to increase the gain by 1.5dB. If the threshold is not exceeded after 32 samples and the DAGC setting is correct, then the noise/speech threshold is recalculated by finding the maximum energy over the last 32 samples (as in B) and multiplying by 1.5 (as in C).

G
Once the noise/speech threshold has been exceeded the filter bank is scanned every 10 milliseconds and the filter data is stored in memory, to form the speech templates, until either 100 samples have been entered or until the energy level drops below the noise/speech threshold for 20 consecutive samples. As described above, if during the data input the A-D converter is overloaded, the AGC setting is decremented by 1 and the data for that filter channel is reprocessed. If during the scan of the 16 filter channels the gain of the DAGC circuit is reduced, the data from all 16 channels is re-input so that <u>all</u> the filter data corresponds to the same AGC setting. The AGC value used is recorded in memory along with the filter data. The AGC setting used at the start of each time slot is taken from the previous time frame, hence the gain can only be reduced (not increased) during the speech processing phase. This is not a problem since at the end of the template period all the template data is normalised to a uniform AGC setting.

H
To ensure that the start of speech was not missed by the speech/noise detector threshold, the 15 time samples prior to speech detection are transferred from the temporary cyclic store to the front of the 'speech' template.

I
If more than 100 samples were processed prior to speech being detected, the noise template is recalculated by analyzing (as in D) the oldest 32 time frames in the temporary cyclic store. If fewer than 100 samples were processed prior to speech being detected, the noise template established in step D is used in the following steps.

J
The minimum gain setting of the AGC over the speech template is then found and both the speech and noise templates are normalised to this setting, which results in both templates containing the values that would have been entered had that gain been used from the start.

K
The normalised noise template is then subtracted from every time frame of the normalised speech template.

L
The maximum energy in the normalised speech template is now found and a new noise/speech threshold calculated - equal to the maximum energy minus 18dB. This new threshold is used to scan the normalised speech template to determine the start and finish points of the speech.

M
The speech template is then truncated to the start and finish points and is either stored in memory (training) or is used for recognition. The following tabular example represents the values stored after measuring the background noise for 320 milliseconds (32 time slots of 10 milliseconds each).

Filter bank number.

	DAGC	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	Real AV energy
	4	210	220	232	245	224	216	167	188	176	234	250	177	134	170	213	209	400
	4	210	218	230	250	220	222	170	190	173	230	255	170	137	172	215	212	409
	4	210	222	234	247	216	225	171	189	178	233	253	171	140	170	214	200	410
	4	213	220	231	251	218	223	166	184	174	230	250	160	133	165	220	216	400
	4	215	217	228	253	220	220	160	186	180	231	254	166	132	164	223	220	409
	4	210	220	232	245	224	216	172	187	177	235	253	160	130	172	214	207	407
	4	210	220	232	245	224	216	167	188	176	234	250	177	134	170	213	209	400
T	4	211	218	230	250	220	222	170	190	173	230	255	170	137	172	215	212	409
I	4	210	222	234	247	216	225	171	189	178	233	253	171	140	170	214	200	410
M	4	213	220	231	251	218	223	166	184	174	230	250	168	133	165	220	216	400
E	4	215	217	228	253	220	220	160	186	180	231	254	166	132	164	223	220	409
	4	210	220	232	245	224	216	172	187	177	235	253	160	130	172	214	207	407
F	4	210	220	232	245	224	216	167	188	176	234	250	177	134	170	213	209	400
R	4	211	218	230	250	220	222	170	190	173	230	255	170	137	172	215	212	409
A	4	210	222	234	247	216	225	171	189	178	233	253	171	140	170	214	200	410
M	4	213	220	231	251	218	223	166	184	174	230	250	168	133	165	220	216	400
E	4	215	217	228	253	220	220	160	186	180	231	254	166	132	164	223	220	409
S	4	210	220	232	245	224	216	172	187	177	235	253	160	130	172	214	207	407
	4	210	220	232	245	224	216	167	188	176	234	250	177	134	170	213	209	400
	4	211	218	230	250	220	222	170	190	173	230	255	170	137	172	215	212	409
	4	210	222	234	247	216	225	171	189	178	233	253	171	140	178	214	200	410
	4	213	220	231	251	218	223	166	184	174	230	250	168	133	165	220	216	408
	4	215	217	228	253	220	220	160	186	180	231	254	166	132	164	223	220	409
	4	210	220	232	245	224	216	172	187	177	235	253	160	130	172	214	207	407
	4	210	222	234	247	216	225	171	189	178	233	253	171	140	170	214	200	410
	4	213	220	231	251	218	223	166	184	174	230	250	168	133	165	220	216	400
	4	215	217	228	253	220	220	160	186	180	231	254	166	132	164	223	220	409
	4	210	220	232	245	224	216	172	187	177	235	253	160	130	172	214	207	407
	4	210	220	232	245	224	216	167	188	176	234	250	177	134	170	213	209	400
	4	211	218	230	250	220	222	170	190	173	230	255	170	137	172	215	212	409
	4	213	220	231	251	218	223	166	184	174	230	250	160	133	165	220	216	400
	4	215	217	228	253	220	220	160	186	180	231	254	166	132	164	223	220	409

Average noise template :-

212 219 231 248 220 220 167 187 176 232 252 169 134 169 217 212

A DAGC value of 4 is equivalent to a 6dB attenuation of the signal going into the A/D, hence to calculate the "real" energy all the filter bank values above would have to be doubled.
Maximum real energy (averaged over all filters) was:- 410
Threshold to be exceeded to start/end template recording:-615

Because the invention's primary application is to voice recognition it has been described with reference to that application. However, as those skilled in the art will be aware, the invention is not only applicable to voice recognition, but is applicable to practically any situation where voice signals are processed for feature extraction.

The speech processor according to the present invention is particularly suitable for use in applications where background noise and variations in the level of that background noise are a problem for known speech processors. One such application is in hands-free telephony, and in particular hands-free telephony involving cellular radio terminals. Such terminals are frequently used in cars, where it is convenient to use speech recognition to provide hands-free call connection and dialling. The problem arises however that wind, road and engine noise fluctuate widely and make accurate recognition of speech difficult. Clearly, if speech recognition for hands-free telephony is to be fully acceptable in this application it is necessary that the recogniser accepts and acts correctly in response to voiced commands in the presence of background noise, without routinely requiring that the commands be repeated.

The improved accuracy of recognition provided by the present invention is of particular advantage in this application.

Claims

1. A method of processing speech in which the start and end points of a speech sample in an input signal are determined by establishing an initial threshold based on a measure of the noise level in said input signal, the initial threshold being used to establish an initial stored speech sample, which initial speech sample is then further processed using a further threshold level which is at a predetermined level beneath a maximum level of said initial speech sample, said further threshold level being used to determine the start and end points.

2. A method as claimed in claim 1, wherein the initial stored speech sample is produced by sampling and storing the input signal to determine when the initial threshold level is exceeded, whereafter an initial reference pattern is stored, the n samples of the input signal which immediately precede the exceeding of said initial threshold level being added to the front end of said initial reference pattern.

3. A method as claimed in claim 1 or claim 2, wherein the determination of said further threshold level is carried out on a signal from which a noise estimate has been subtracted.

4. A method as claimed in any one of the preceding claims, wherein a normalisation step is carried out prior to the determination of said further threshold level.

5. A method as claimed in any one of the preceding claims, wherein a speech template is produced and stored, the speech template being truncated to the start and end points determined using said further threshold.

6. A method of processing speech comprising the steps of: determining a speech energy threshold for an input channel;

   forming a background noise template for the input channel;

   sampling the input channel and storing the samples into a first store;

   detecting speech by determining when a sample exceeds the speech threshold, and storing the sample and subsequent samples into a second store;

   detecting the end of the speech by determining when a predefined number of the subsequent samples drop below the threshold;

   forming a first speech template by augmenting the samples from the second store with a predefined number of samples from the first store; and

   forming a second speech template by subtracting the background noise template from the first speech template.

7. A method of processing speech as claimed in claim 6 in which the step of determining the speech threshold for the input channel comprises averaging the background noise of the input channel over a first time period and setting the speech threshold to be H times greater than the average background noise.

8. A method of processing speech as claimed in claim 6 or claim 7 in which the background noise template comprises a plurality of consecutive samples of the input channel, the samples being taken when no speech is present.

9. A method of processing speech as claimed in any of claims 6 to 8 in which the first store is a cyclic store.

10. A method of processing speech as claimed in claim 9 wherein the cyclic store comprises 100 addressable store locations.

11. A method of processing speech as claimed in any of claims 6 to 10 in which the second store is a random access memory.

12. A method of processing speech as claimed in any of claims 6 to 11 in which he step of forming the first speech template comprises augmenting the samples from the second store by adding a predefined number of samples from the first store, the predefined number of samples being those which directly preceded the first sample which exceeded the speech threshold, in front of the stored samples in the second store.

13. A method of processing speech as claimed in any of claims 6 to 12 further comprising a step of checking a gain factor of the system, whereby, if the speech threshold is not exceeded within a predefined number of samples and if all the samples in the first store are more than H times smaller than the speech threshold, the gain factor is increased by H times.

14. A method of processing speech as claimed in any of claims 6 to 13 wherein if the speech threshold is not exceeded within a predefined number of samples and not all the samples in the first store are more than H times smaller than the speech threshold, the speech threshold is re-calculated by finding the maximum energy sample over a preceding predefined number of samples and setting the new speech threshold to be H times greater than the maximum energy sample.

15. A method of processing speech as claimed in any of claims 6 to 14 in which, if more than a predefined number of samples were processed prior to speech being detected, the background noise template is re determined prior to the formation of the second speech template.

16. A method of processing speech as claimed in any of claims 6 to 15 further comprising a step of forming a third speech template from all samples in the second speech template which exceed a second threshold, the second threshold being set at a predetermined level below the maximum energy level of the second speech template.

17. A speech processing apparatus which processes speech according to the method of any one of the preceding claims.

18. A speech processing apparatus as claimed in claim 17 configured as a speech recognizer.

19. A speech processing apparatus in which the start and end points of a speech sample in an input signal are determined according to a method in which:

(i) a measure of noise level prevailing in said input signal is determined;

(ii) a signal threshold level T, greater than the determined level, is established;

(iii) the input signal is sampled and stored to determine when threshold level T is exceeded, whereafter an initial reference pattern is stored;

(iv) the n samples of the input signal immediately preceding the exceeding of said threshold level T are added to the front end of said initial reference pattern;

(v) a new threshold level R is derived from the pattern produced in step (iv) and this new threshold is used to scan a stored signal derived from said initial reference pattern to determine the start and finish points of speech.

20. A telephony terminal including a speech processor as claimed in claim 17, claim 18 or claim 19.

Fig.1.

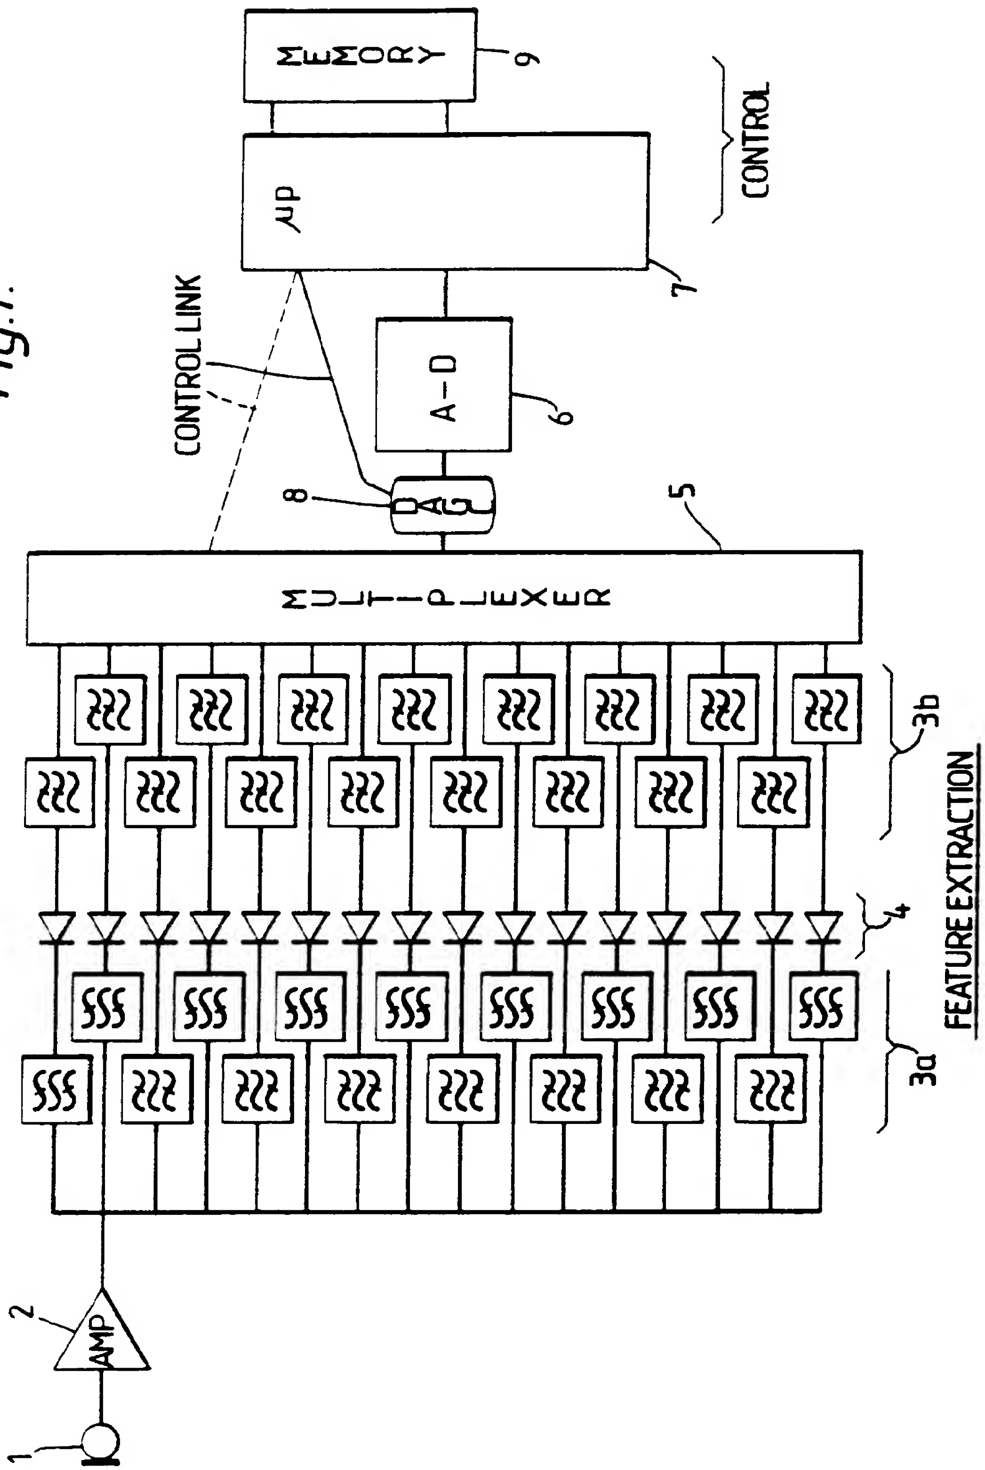
MEMORY
9
μP
CONTROL
7
CONTROL LINK
A-D
6
8
AGC
5
MULTIPLEXER
3b
4
3a
FEATURE EXTRACTION
2
AMP
1

European Patent Office

EUROPEAN SEARCH REPORT

Application Number

EP 93 20 3407

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
X	THE BELL SYSTEM TECHNICAL JOURNAL, vol. 54, no. 2, February 1975, pages 297-315, American Telephone and Telegraph Co., Murray Hill, New Jersey, US; L.R. RABINER et al.: "An algorithm for determining the endpoints of isolated utterances" * Figure 9 * ---	1,2,6	G 10 L 3/00
A	ICASSP 83, IEEE International Conference on Acoustics, Speech and Signal Processing, April 14-16, 1983, vol. 1, pages 312-315, IEEE, New York, US; P. FONSALE: "Connected-word recognition system using speaker-independent phonetic features" * Page 313, "preprocessing" * ---	2,19	
A	ICASSP 85, IEEE International Conference on Acoustics, Speech, and Signal Processing, March 26-29, 1985, Tampa, US, vol. 4, page 1838, IEEE, New York, US; M-I.CHUNG et al.: "Word boundary detection and speech recognition of noisy speech by means of iterative noise cancellation techniques" * Abstract * ---	3	TECHNICAL FIELDS SEARCHED (Int. Cl.5)
			G 10 L 3/00
A	WO-A-8 603 047 (AMERICAN TELEPHONE AND TELEGRAPH CO.) * Figure 1 * ---	4	
A	GB-A-2 075 312 (STORAGE TECHNOLOGY CORP.) * Page 1, lines 104-118 * --- -/-	4	

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
THE HAGUE	21-01-1994	ARMSPACH J F A

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P0401)

European Patent Office

EUROPEAN SEARCH REPORT

Application Number

EP 93 20 3407

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
A	US-A-4 410 763 (STRAWCZYNSKI et al.) * Claim 1 * ---	7	
A	EP-A-0 025 684 (INTERSTATE ELECTRONICS CORP.) * Abstract * ---	9	
A	US-A-4 052 568 (JANKOWSKI) * Column 4, line 25 - column 5, line 2 * ---	13,14, 15	
A	AT&T BELL LABORATORIES TECHNICAL JOURNAL, vol. 63, no. 3, March 1984, pages 479-497, Murray Hill, New Jersey, US; J.G. WILPON et al.: "An improved word-detection algorithm for telephone-quality speech incorporating both syntactic and semantic constraints" * Paragraph II: "Review of bottom-up endpoint detector" * ---	17,18, 20	TECHNICAL FIELDS SEARCHED (Int. Cl.5)
A	FR-A-2 451 680 (SOUMAGNE) * Claim 1 * ---	1,6,19	
A	FR-A-2 304 135 (THRESHOLD TECHNOLOGY INC.) * Claim 1 * ---	1,6,11, 19	
P,X	US-A-4 627 091 (FEDELE) * Abstract * --- -/-	1,2,6, 11,19	

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
THE HAGUE	21-01-1994	ARMSPACH J F A

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P0401)

European Patent Office

EUROPEAN SEARCH REPORT

Application Number

EP 93 20 3407

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
E	EP-A-0 237 934 (K.K. TOSHIBA) * Claims 8-13 * -----	1,6,7, 19	

TECHNICAL FIELDS SEARCHED (Int. Cl.5)

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
THE HAGUE	21-01-1994	ARMSPACH J F A

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons
& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P0401)